REMARKS

Claims 5, 11 and 17 have been canceled without prejudice or disclaimer of the subject matter thereof. Applicants reserve the right to pursue the subject matter of the canceled claims in subsequently filed continuing applications.

Claims 1, 6 - 7, 12 - 13 and 18 have been amended.

Claims 1 - 4, 6 - 10, 12 - 16 and 18 are present in the subject application.

In the Office Action dated January 2, 2008, the Examiner has rejected claims 1 - 18 under 35 U.S.C. §103(a). Favorable reconsideration of the subject application is respectfully requested in view of the following remarks.

The Examiner has rejected claims 1 - 2, 4 - 8, 10 - 14 and 16 - 18 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 5,625,815 (Maier et al.) in view of U.S. Patent No. 6,289,334 (Reiner et al.). This rejection is considered moot with respect to canceled claims 5, 11 and 17.

Briefly, the present invention is directed toward a database management system installed in a data processing system. The database management system manages a database having partitions for storing table data based on a partitioning schema, in which each partition has an associated partition identifier, and in which the database has database catalog information associated therewith. A partition identifier is identified in accordance with the partitioning schema, and the partition identifier is selected based on the contents of the query and the database catalog information. The query is executed against the identified partition. The technique improves the execution of queries while minimizing the consumption of network resources. For example, the technique eliminates the processing costs associated with directing

the query to the appropriate partition, and transferring data from the data partition to a corresponding remote coordinator handling the query statement.

The Examiner takes the position that the Maier et al. patent discloses the claimed subject matter, except for executing the query against the specific partition. The Examiner further alleges that the Reiner et al. patent discloses this feature, and that it would have been obvious to combine the Maier et al. and Reiner et al. patents to attain the claimed invention.

This rejection is respectfully traversed since the Maier et al. and Reiner et al. patents do not disclose, teach or suggest determining the specific partition containing the database table portion with the desired data, identifying a partition identifier associated with the specific partition and executing the query against the specific partition as recited in the claims. However, in order to expedite prosecution of the subject application, independent claims 1, 7 and 13 have been amended to further clarify these features and recite the features of: each networked partition being contained within a corresponding network processing node; a plurality of agent modules each associated with a corresponding networked partition to establish a physical connection with, execute a client query against and retrieve desired data from that networked partition; providing a client query for the database at a network processing node in accordance with a user request to retrieve the desired data; analyzing the contents of the client query and the retrieved database catalog information to determine a specific partition from among the plurality of networked partitions containing the database table portion with the desired data satisfying the client query and identifying a partition identifier associated with the specific partition; and executing the client query against the specific partition by determining the associated agent module for the specific partition based on the partition identifier and directing the determined

agent module to execute the client query against, and retrieve the desired data satisfying the client query from, the specific partition.

The Maier et al. patent does not disclose, teach or suggest these features. Rather, the Maier et al. patent is directed toward a database computer system including a memory, residing in a plurality of interconnected computer nodes, for storing database tables. Each database table has a plurality of columns, a primary key index based on a specified subset of the columns, and an associated table schema. At least a subset of the database tables are partitioned into a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from the other partitions. A transaction manager generates and stores an audit trail, each audit entry denoting a database table record event, such as an addition, deletion or alteration of a specified database table record in a specified one of said database tables (e.g., See Abstract).

Four online data definition procedures allow the structure of a database table to be altered while the database table remains available to execution of transactions, with minimal impact of the availability of the database table for transaction execution. The four online data definition procedures are a move partition procedure, a split partition procedure, a move partition boundary procedure, and a create new index procedure (e.g., See Abstract; Column 2, lines 39 - 46; Column 7, lines 61 - 63; Column 9, lines 14 - 17; Column 10, lines 47 - 50; and Column 12, lines 19 - 22). Each of these online procedures has three or four phases of execution and begins with a user or operator entering a DDL statement specifying an alteration in the schema of a specified object, typically either a database table, an index or a partition (e.g., See Column 6, lines 9 - 13).

In a first phase, records of a table partition or the entire table are accessed using read only access, so as to generate a new partition, move records between two partitions, or to create a new

index (e.g., See Abstract; Column 2, lines 47 - 57; Column 8, lines 15 - 29; Column 9, lines 41 - 50; Column 11, lines 10 - 20; and Column 12, lines 39 - 42). In a second phase, audit trail entries are used to clean up the data structures created during the first phase (e.g., See Abstract; Column 2, lines 58 - 62; Column 8, lines 30 - 50; Column 9, line 51 to Column 10, line 3; Column 11, lines 20 - 42; and Column 13, lines 12 - 22). In a third phase, access to the database table is briefly locked while audit trail entries created after the second phase are used to make final changes to the data structures created during the first phase, and while the database table schema is updated to reflect the changes to the database table produced (e.g., See Abstract; Column 2, lines 62 - 67; Column 8, line 51 to Column 9, line 10; Column 10, lines 4 - 40; Column 11, line 43 to Column 12, line 12; and Column 14, lines 11 - 30). In a fourth phase, used by the move partition boundary and split partition procedures, records in a database partition that are inconsistent with the modified database schema are deleted as a background operation while use of the database table by transactions resumes (e.g., See Column 3, lines 1 - 5).

Thus, the Maier et al. patent is directed toward modifying a database table's structure in accordance with a user command specifying the object (e.g., partition), while maintaining access to database tables during the modification. There is no disclosure, teaching or suggestion of: networked processing nodes with agent modules each corresponding to a networked partition to establish a physical connection, execute a query and retrieve desired data from that networked partition; determining the specific partition containing the desired data satisfying a query based on the query contents and database organization or catalog information relating to the database table containing the desired data; and executing the query against that specific partition by

determining the associated agent module for the specific partition based on the partition identifier and directing the determined agent module to execute the client query and retrieve the desired data from the specific partition as recited in the independent claims.

The Examiner indicates in the Office Action that the feature of determining a specific partition is inherent in the Maier et al. patent and well known in the art. This position is respectfully traversed. The Maier et al. patent does not disclose, teach or suggest the claimed manner of query execution. As the Examiner is well aware, inherency may not be established by probabilities or possibilities, or from the mere fact that a specific occurrence may result from a given set of circumstances. See M.P.E.P. §2112. The circumstance that a certain result may occur or be present in the prior art is insufficient to establish inherency of that result. Id. The Examiner must provide a basis in fact and/or technical reasoning to support the determination that the inherent characteristic necessarily flows from the applied prior art. Id.

The Reiner et al. patent does not compensate for the deficiencies of the Maier et al. patent. Rather, the Reiner et al. patent is directed toward intercepting database queries prior to reception by a database management system (DBMS) and decomposing selected queries to generate multiple sub-queries for application, in parallel, to the DBMS (e.g., See Abstract; Column 2, line 65 to Column 3, line 6; and Column 7, line 64 to Column 8, line 51).

Thus, the Reiner et al. patent intercepts an initial query and produces a plurality of sub-queries for execution against database partitions (e.g., See Fig. 3B). If a query is not decomposable, it is transferred to the DBMS without being decomposed and processed in a conventional manner by a server process applying the query to each of the partitions in a non-parallel fashion (e.g., See Figs. 3A - 3B; and Column 7, line 64 to Column 8, line 1).

Accordingly, there is no disclosure, teaching or suggestion of: networked processing nodes with agent modules each corresponding to a networked partition to establish a physical connection, execute the client query (the initial query based on a user request) and retrieve desired data from that networked partition; determining the specific partition containing the desired data satisfying the client query based on the client query contents and database organization or catalog information relating to the database table containing the desired data; and executing the client query against that specific partition by determining the associated agent module for the specific partition based on the partition identifier and directing the determined agent module to execute the client query and retrieve the desired data from the specific partition as recited in the independent claims.

In fact, the Reiner et al. patent discloses decomposing the initial query into sub-queries in order to take advantage of processing the query in parallel (e.g., See Column 2, lines 31 - 36; Column 3, lines 2 - 3; and Column 8, lines 46 - 50), thereby teaching away from applying the initial query to the database partitions. Further, the Maier et al. patent is directed toward modifying database table structures in accordance with user commands specifying the object or partition to modify as described above. Thus, there is no apparent reason to combine query processing with these types of processes.

Since the Maier et al. and Reiner et al. patents do not disclose, teach or suggest, either alone or in combination, the features recited in independent claims 1, 7 and 13 as discussed above, these claims are considered to be in condition for allowance.

Claims 2, 4, 6, 8, 10, 12, 14, 16 and 18 depend, either directly or indirectly, from independent claims 1, 7 or 13 and, therefore, include all the limitations of their parent claims.

Claims 6, 12 and 18 have been amended for consistency with their amended parent claims. The dependent claims are considered to be in condition for allowance for substantially the same reasons discussed above in relation to their parent claims and for further limitations recited in the dependent claims.

The Examiner has rejected claims 3, 9 and 15 under 35 U.S.C. §103(a) as being unpatentable over the Maier et al. and Reiner et al. patents, and further in view of U.S. Patent Application Publication No. 2003/0233347 (Weinberg et al.). Briefly, the present invention is directed toward a database management system installed in a data processing system as described above.

The Examiner takes the position that the Maier et al. and Reiner et al. patents disclose the claimed subject matter, except for building a subset of the database catalog. The Examiner further alleges that the Weinberg et al. publication discloses this feature, and that it would have been obvious to combine the Maier et al. and Reiner et al. patents with the Weinberg et al. publication to attain the claimed invention.

This rejection is respectfully traversed. Initially, claims 3, 9 and 15 depend, either directly or indirectly, from independent claims 1, 7 or 13 and, therefore, include all the limitations of their parent claims. As discussed above, the Maier et al. and Reiner et al. patents do not disclose, teach or suggest the features of: networked processing nodes with agent modules each corresponding to a networked partition to establish a physical connection, execute a client query and retrieve desired data from that networked partition; determining the specific partition containing the desired data satisfying a query based on the query contents and database organization or catalog information relating to the database table containing the desired data; and

executing the query against that specific partition by determining the associated agent module for the specific partition based on the partition identifier and directing the determined agent module to execute the client query and retrieve the desired data from the specific partition as recited in the claims.

The Weinberg et al. publication does not compensate for the deficiencies of the Maier et al. and Reiner et al. patents. Rather, the Weinberg et al. publication is directed toward structuring, storing and retrieving data in database systems utilizing relational objects or qualifiers, and is merely utilized by the Examiner for an alleged teaching of a subset of catalog data. However, the catalog data relied on by the Examiner refers to items of manufacture or products (e.g., See Paragraph 0046), as opposed to database catalog information as recited in the claims.

Since the Maier et al. and Reiner et al. patents and Weinberg et al. publication do not disclose, teach or suggest, either alone or in combination, the features recited in claims 3, 9 and 15 as discussed above, these claims are considered to be in condition for allowance.

The application, having been shown to overcome the issues raised in the Office Action, is considered to be in condition for allowance and a Notice of Allowance is earnestly solicited.

Respectfully submitted,

Stuart B. Shapiro
Registration No. 40,169

EDELL, SHAPIRO & FINNAN, LLC
1901 Research Blvd., Suite 400
Rockville, Maryland 20850-3164
(301) 424-3640
Delivered: 02-01-08

16